

HP StorageWorks

Clustered File System 3.0

setup guide



393721-001

Part number: 393721-001
First edition: May 2005



Legal and notice information

© Copyright 1999-2005 PolyServe, Inc.

Portions © 2005 Hewlett-Packard Development Company, L.P.

Neither PolyServe, Inc. nor Hewlett-Packard Company makes any warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Neither PolyServe nor Hewlett-Packard shall be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Neither PolyServe nor HP shall be liable for technical or editorial errors or omissions contained herein.

The software this document describes is PolyServe confidential and proprietary.

PolyServe and the PolyServe logo are trademarks of PolyServe, Inc.

PolyServe Matrix Server contains software covered by the following copyrights and subject to the licenses included in the file *thirdpartylicense.pdf*, which is included in the PolyServe Matrix Server distribution.

Copyright © 1999-2004, The Apache Software Foundation.

Copyright © 1992, 1993 Simmule Turner and Rich Salz. All rights reserved.

Copyright © 2000, Emulex Corporation. All rights reserved.

Copyright © 1994-2004, Sun Microsystems, Inc. All rights reserved.

Copyright © 1999, 2000, 2001 Lev Walkin <vlm@spelio.net.ru>. All rights reserved.

Copyright © 1996, 1997, 1998, 1999, 2000, 2001 Institut National de Recherche en Informatique et en Automatique (INRIA). All rights reserved.

Copyright © 2001 QLogic Corp. All rights reserved.

Copyright © 1993-2001 Spread Concepts LLC. All rights reserved.

Copyright © 2001 Daniel Barbalace. All rights reserved.

Copyright © 2003 Storage Networking Industry Association. All rights reserved.

Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Linux is a U.S. registered trademark of Linus Torvalds

Microsoft®, Windows®, Windows NT®, and Windows XP® are U.S. registered trademarks of Microsoft Corporation.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

UNIX® is a registered trademark of The Open Group.

Clustered File System 3.0 setup guide

Contents

HP Technical Support

HP Storage Web Site	v
HP NAS Services Web Site	v

1 Configuration Information

Configuration Limits	1
Supported HBA Drivers	2
Software Image Overview.....	2
NTP Requirements.....	3
Management Console Requirements	3

2 Setting Up Physical Connectivity

Setting Up the HP StorageWorks DL380-SL Clustered Gateway ..	4
Cluster Nodes	4
Client Network Connections	4
Cluster SAN Configuration Guidelines	6

3 Setup Procedure

Initial Pre-Planning Steps	9
Setup Checklist	10
1. Power on Each Node and Login.....	12
2. Configure the Storage Array.....	14
3. Configure FC Switches for the Cluster	14
4. Create LUNs or Disk Partitions for Membership Partitions .	16
5. Configure Node Settings.....	16
6. Verify the SAN Configuration	19
7. Run the mxcheck Utility	19
8. Run the mxconfig Utility.....	20
9. Start HP Clustered File System on One Server	38

10. Start the Management Console	38
11. Add the Remaining Servers to the Cluster	39
12. Start HP Clustered File System on the Remaining Servers ..	40
13. Configuring HP CFS for Public versus Private Network ..	40
14. Set Up NFS file services.....	41
15. Save System State	41

4 Post Setup Procedures

Verify the Fencing Configuration	43
Install the Management Console on Other Hosts	44
Management Console for Linux Systems	44
Management Console for Windows Systems.....	45
Disaster Recovery.....	46
Re-Image the Cluster Node.....	47
Restore the Node State.....	47
CIFS File Serving	48
Quota Management.....	48

5 Other Procedures

Add a New Server to a Cluster	50
Upgrade the License File.....	51
Change the Fencing Method.....	52
Uninstall the Management Console	52
Set HP Clustered File System Parameter for FalconStor	53

A Configure the Cluster from the Management Console

Cluster Configuration	54
Required Information	54
Configuration Procedure.....	55
Test the Fencing Configuration	68

Index

HP Technical Support

Telephone numbers for worldwide technical support are listed on the following HP web site: <http://www.hp.com/support>. From this web site, select the country of origin. For example, the North American technical support number is 800-633-3600.

NOTE: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Web Site

The HP web site has the latest information on this product, as well as the latest drivers. Access the storage site at:

<http://www.hp.com/country/us/eng/prodserv/storage.html>. From this web site, select the appropriate product or solution.

HP NAS Services Web Site

The HP NAS Services site allows you to choose from convenient HP Care Pack Services packages or implement a custom support solution delivered by HP ProLiant Storage Server specialists and/or our certified service partners. For more information see us at

http://www.hp.com/hps/storage/ns_nas.html.

Configuration Information

HP is continually expanding its supported hardware and operating system configurations. For the latest information, check the *HP StorageWorks Clustered File System Hardware and Software Compatibility Guide* on the HP web site:

1. Go to <http://www.hp.com/go/nas>.
2. Select the link for the HP StorageWorks Enterprise File Services DL380-SL Clustered Gateway product.
3. Select the “Technical documentation” link.
4. Select the link to the *HP File Serving Cluster Hardware and Software Compatibility Guide*.

Configuration Limits

The following hardware can be used in an HP Clustered File System configuration:

Hardware	Configuration Limit
Servers	Two to 16 servers.
Network Interface Cards	Up to four network interfaces per server.

Hardware	Configuration Limit
FibreChannel Host Bus Adapters	Four FC ports per server can be connected to the cluster SAN. Other FC ports can be connected to non-cluster SANs.
FibreChannel Switches	Two levels of cascading switches.
FibreChannel Storage Subsystems	Up to 508 LUNs.

Supported HBA Drivers

HP supports the driver that comes pre-installed on the HP StorageWorks Clustered File System.

HP will continue to evaluate newer drivers with HP Clustered File System as they become available from the vendors, and will publish updated support information on the HP web site. Check the *HP StorageWorks Clustered File System Hardware and Software Compatibility Guide* before upgrading to a newer driver from your HBA vendor. Refer to page 1 for location information.

NOTE: QLogic and Emulex cards should not be placed in the same fabric zone unless both HBA manufacturers have approved this configuration.

Software Image Overview

HP StorageWorks DL380-SL Clustered Gateway servers come with all required software pre-installed. The major components include:

- SuSE Linux Enterprise Server 9
- HP EFS Clustered Software
- HP EFS Clustered Volume Manager (optional license required)
- HP Insight Manager Home Page & Agents
- HP Quick Configuration Management Tools

Each server has two SCSI drives. The operating system and associated software is mirrored across the two disks using the built-in HP Smart Array RAID controller.

A quick restore DVD is included with each server, should you need to re-image the server. See "Disaster Recovery" on page 15 for details.

NTP Requirements

NOTE: To ensure that file times are consistent across the cluster, it is important that all cluster servers operate with synchronized time-of-day clocks. An NTP server is one commonly used mechanism for synchronizing system clocks.

Management Console Requirements

Servers running the Management Console must have a windowing environment installed. The Management Console requires that the display be set to a minimum of 256 colors.

Setting Up Physical Connectivity

This chapter describes how to physically configure a new HP StorageWorks Enterprise File Services (EFS) Clustered Gateway.

Setting Up the HP StorageWorks DL380-SL Clustered Gateway

Before attempting to power on the nodes on an HP StorageWorks EFS Clustered Gateway, there is required hardware setup. The HP support specialist who installs your system performs most of these tasks. An overview of the required hardware setup is provided below.

Cluster Nodes

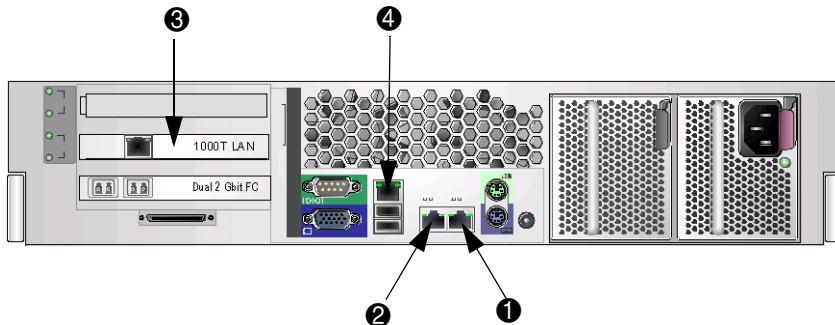
Some HP StorageWorks EFS Clustered Gateway solutions come pre-racked while others will require the equipment to be racked after it arrives.

Servers come with the required hardware for installing them in a standard HP rack.

Client Network Connections

It is important that you use the IP Network connections correctly. Failure to follow these guidelines could result in a poorly performing system or a non-functional system. This section explains which NIC ports to use.

The illustration below identifies the network connections.



Item	Description
1	Integrated NIC 1 — Intra-cluster communication (private network, administration traffic)
2	Integrated NIC 2 — Client network connection (public network, NFS traffic)
3	PCI Card NIC — Client network connection (public network, NFS traffic)
4	iLO NIC — Used for server fencing (private network, administration traffic)

Networks used with HP Clustered File System should be configured as follows:

- Each network interface card (NIC) must be on a separate physical network. For example, you could use either a separate Ethernet switch for each network or a managed switch that allows network traffic to be isolated.
- Each NIC must be on a separate logical IP network.
- The network topology should be symmetrical; each server in the cluster should be connected to the same set of networks.

The following flags must be set on each network interface: Broadcast, Multicast, Up, Running.

Client Network Connections

Each server has two Gigabit Ethernet connections that are to be used for data (NFS) connectivity to the public network. The internal NIC labeled 2 on the back of the DL380 and the single port NIC card in the middle PCI slot should be used for public network traffic.

NOTE: It is important to configure the software so it does not use the public network connections for intra-cluster network traffic. Refer to “13. Configuring HP CFS for Public versus Private Network” on page 40 for details.

Intra-Cluster Network Connections

The cluster servers must be interconnected using a private network. This interconnect is critical to keep the servers in the cluster synchronized. Overall cluster status, health and intra-server communication happens on this network. The internal NIC labeled 1 on the back of the DL380 should be used for this network traffic. All such intra-client traffic must use a private network. Either a dedicated switch or a switch with VLAN capabilities should be used. A dedicated network switch is included for this purpose with some HP StorageWorks EFS Clustered Gateway configurations.

iLO Card Network Connections

The DL380 G4 servers include a built-in remote management card called the Integrated Lights Out (iLO) card. The iLO plays an integral part in the HP StorageWorks EFS Clustered Gateway solution. It is used to restart a server in the cluster if the server is not adhering to the rules of cluster membership. The iLO card has a network connection. All iLO cards should be part of the same intra-cluster network described above.

Cluster SAN Configuration Guidelines

DL380 G4 cluster servers have a Fibre Channel HBA installed on the bottom PCI slot. The placement of the HBA is important as there are several different speed PCI busses in a DL380 G4 server.

HP recommends that the HBA not be moved to another slot to maintain the best performance. The HBA is a dual-channel model: one channel is used for the primary path to the SAN subsystem; the other channel is used as a secondary or redundant path to the storage. With some storage arrays, the second channel could be also be used as a load-balanced path, as well as a failover path.

Following are guidelines for configuring the cluster SAN to be used with HP Clustered File System:

- The FibreChannel fabric used for an HP Clustered File System cluster can be shared with other HP Clustered File Systems matrices or with non-cluster servers and devices. To protect the cluster SAN devices from other devices on the fabric, particularly devices controlled by Windows platforms, we recommend that you create a fabric zone for the cluster.
- The FC switches managed by the cluster must be from the same vendor. Switches that will not be managed by HP Clustered File System can be from other vendors only if all of the vendors have approved the configuration.
- Multiported disks are supported on the cluster SAN. The fabric handles the multiple ports transparently to the operating system.
- To eliminate single points of failure, configure the SAN to use multipath I/O. See the HP StorageWorks Clustered File System Administration Guide for some sample cluster configurations that include this feature.
- Disk Power Management must be disabled on any Windows systems that have access to the SAN storage.

NOTE: You can attach a cluster server to other SANs that are not under HP Clustered File System control. These SANs can include tape drives and non-PSFS filesystems. You can use Linux utilities or other applications to access these devices and filesystems.

The design of the SAN storage is not covered in this document. A variety of storage arrays and fabric switches are available.

Contact your HP sales representative for assistance in planning and implementing the storage component of the overall solution.

3

Setup Procedure

This chapter describes how to configure a new HP StorageWorks Clustered File System.

CAUTION: Before starting this procedure, review the configuration and hardware requirements in Chapter 1.

Initial Pre-Planning Steps

Complete the HP Pre-Planning Worksheet with your HP service expert. This process is part of the HP installation service associated with this product. Your HP representative provides this worksheet. A copy is also available on the documentation CD with this manual.

It is important to fill out the worksheet before you start configuring the cluster nodes to ensure that you have the information required to correctly configure your solution.

The following items are required for initial node setup:

- Network parameters (IP addresses, subnet masks, etc)
 - The two data NICs on each node
 - The intra-node heartbeat NIC
 - The iLO card

- Host names and DNS servers
- NTP Server
- SAN Storage I Cluster Node Startup

Setup Checklist

There is some setup that must be done on each server in your cluster. The setup procedure consists of the following steps. Each step is described in detail following this checklist. If you will be using a client host outside the cluster to administer your cluster, you will also need to install the Management Console on that host. The Management Console is available for either Linux or Microsoft® Windows®.

Action	Description	
Power on each node.	<ul style="list-style-type: none"> Configure the iLO card at boot — This is required for the server based fencing part of the solution. Configure the display subsystem — A keyboard and display are required to configure each node. 	<input type="checkbox"/>
Obtain HBA WWID strings.	<ul style="list-style-type: none"> The HBA WWID strings are needed to correctly present LUNs to the cluster nodes. Use any of the following tools to obtain the HBA WWID strings: <ul style="list-style-type: none"> – HP Clustered Gateway Control Panel – /opt/hpcfsmgmt/bin/nasconfig command line tool – Any equivalent Linux tool 	<input type="checkbox"/>
Configure the storage array.	Set up the storage array as described in the array product documentation.	<input type="checkbox"/>
Configure FC switches.	The FibreChannel switches that will be under cluster control need to be configured to enable cluster operations across the SAN.	<input type="checkbox"/>

Action	Description	
Create LUNs or disk partitions for use as membership partitions.	HP Clustered File System uses membership partitions to control access to the SAN. If your disk array software provides the capability to create LUNs, we recommend that you create three LUNs for the membership partitions.	<input type="checkbox"/>
Configure node settings.	Use any of the following tools: <ul style="list-style-type: none"> • HP Clustered Gateway Control Panel • <code>/opt/hpcfsmgmt/bin/nasconfig</code> command line tool • Any equivalent Linux tool 	<input type="checkbox"/>
Verify the SAN configuration.	Verify that the SAN devices are visible from the servers and are configured correctly.	<input type="checkbox"/>
Run the mxcheck utility.	Run the mxcheck utility on each server to verify that the server's configuration meets the requirements for HP Clustered File System.	<input type="checkbox"/>
Run the mxconfig utility or the HP Clustered File System Management Console (mxconsole).	The mxconfig utility, as well as the configuration mode of the mxconsole tool, both perform some initial configuration of HP Clustered File System. Run one of these tools on one server and then export the resulting configuration to the other servers.	<input type="checkbox"/>
Start HP Clustered File System on one server.	Do not start HP Clustered File System on the other servers.	<input type="checkbox"/>
Start the Management Console.	The Management Console is used to configure and manage the cluster. Start the Management Console and connect to the server where you started HP Clustered File System.	<input type="checkbox"/>

Action	Description	
Add the remaining servers.	Add the remaining servers to the cluster one-at-a-time.	<input type="checkbox"/>
Start HP Clustered File System on the remaining servers.	After HP Clustered File System is running on all servers, the installation is complete.	<input type="checkbox"/>
Configure HP CFS for public versus private network.	Use mxconsole to make sure the public network is not used as a primary intra-cluster communication channel.	<input type="checkbox"/>
Set up NFS file services.	After the cluster is up and running, file systems and NFS exports can be configured for the cluster from any node.	<input type="checkbox"/>
Save system state.	Optional step. It is important to create a backup of each node's state should it become necessary to rebuilt it.	<input type="checkbox"/>
Start the MSA Array ASM daemon.	Optional step. If MSA storage is being used, the Array Control Utility can be run directly from any node in the cluster.	<input type="checkbox"/>
Perform post setup tasks	Refer to Chapter 4 on page 43.	<input type="checkbox"/>

1. Power on Each Node and Login

One by one, power on each node in the cluster. You will need information from the HP Pre-Planning Worksheet for this task. Complete the following steps for each node:

1. Attach a keyboard, mouse, and monitor to the system.
2. Power on the system.
3. Configure the iLO card.

The iLO card is an integral part of the solution. The HP Clustered File System software uses the capabilities of the card to do server-based fencing. Configure the card with an IP address, user name, and password that is used later.

- a. As the system is starting, configure the iLO card.
- b. When the Integrated Lights Out line is displayed, press F8.
- c. If prompted for a password, locate the default password for the iLO on the tag attached to the front of the HP StorageWorks DL380-SL Clustered Gateway.
- d. Add a user called HPCFS and a password. Refer to the Integrated Lights Out documentation included with the HP StorageWorks DL380-SL Clustered Gateway.
- e. Set up networking using either a static IP address or DHCP using long-term leases, as appropriate for your network.
- f. Choose File > Exit to continue starting the system.

4. Configure the display subsystem.

When the system first starts, it allows you to configure a display subsystem (monitor type, resolution, etc). After the system is up and running, it can be managed without a keyboard and monitor attached. However, initial system configuration requires you to complete this process. In most cases, the system auto-detects and suggests appropriate settings. The defaults will probably allow you to configure the system. However, HP recommends that you ensure a resolution of at least 1024x768 and a color depth of at least 16 bits.

NOTE: You must have a monitor and keyboard to configure each node.

5. Log in to the system.

When the login screen is displayed, log in as the “root” user. Enter the default password “hpinvent.”

6. Obtain HBA WWID (World Wide ID's)

The Fibre Channel WWID string for each of port (there are two ports on each HBA) will be required to correctly configure the storage subsystem. Use the HP Clustered Gateway Control Panel, the `/opt/hpcfsmgmt/bin/nasconfig` command line tool, or any equivalent linux tool to obtain HBA WWID strings.

The HP tools will execute the script `/opt/hpcfsmgmt/bin/cfshbainfo` to obtain this information. The script looks at the files in `/proc/scsi/q1a2xxx` for the required information. The HBA driver must be loaded for the file entered in `/proc/scsi` to exist (the `cfshabinfo` script takes the required steps to load and unload drivers as required).

2. Configure the Storage Array

Perform the initial configuration of the storage array as described in the array product documentation.

3. Configure FC Switches for the Cluster

When certain problems occur on a server (for example, hardware problems or the loss of cluster network communications), and the server ceases to effectively coordinate and communicate with other servers in the cluster, HP Clustered File System must remove the server's access to filesystems to preserve data integrity. This step is called *fencing*.

When you configure the cluster later in this procedure, you can select the fencing method that you want to use:

- FibreChannel Switch-based fencing (also called "fabric fencing"). When a server needs to be fenced, HP Clustered File System disables the server's access in the FibreChannel fabric.
- Web Management-Based Fencing via Server Reset/Shutdown (also called "server-based fencing"). HP Clustered File System uses the iLO card on the HP ProLiant server to remove its access to PSFS filesystems.

For sites using fabric fencing, the FibreChannel switches must be configured as described below to enable cluster operations across the SAN. Configuring the FibreChannel switches is not required for sites using server-based fencing. However, if the switches are configured as described here, certain HP Clustered File System commands can return more information.

To configure FibreChannel switches, complete the following tasks:

- Enable server access to the SAN. Each server that will be in the cluster must be able to see the disks in the SAN. You may need to enable server ports on the FC switches or to change the zoning configuration to give servers the necessary access to the SAN.
- Modify the switches SNMP setup. Make the following changes:
 - Enable access to the SNMP agent from each server that will be in the cluster.
 - Set the SNMP community string to the desired value. The community string must be the same for all FibreChannel switches. If you are using fabric fencing, ensure that the community string has write privileges. For Web Management-Based Fencing, you can optionally configure a read-only community string.
 - On Brocade switches only, run the **snmpMibCapSet** command on the switch. Change the **famib** setting to **yes** and accept the default values for the other settings.
- For Cisco and McDATA switches, each server should be placed in a separate zone with its storage.

Typically these tasks can be performed from the FC switch or from third-party applications. Refer to your FC switch or application documentation for more information.

4. Create LUNs or Disk Partitions for Membership Partitions

HP Clustered File System uses a set of membership partitions to control access to the SAN and to store the device naming database, which includes the global device names that HP Clustered File System assigns to the SAN disks placed under its control.

HP Clustered File System can use either one or three membership partitions. To ensure that a membership partition is always available, we strongly recommend that you use three membership partitions.

If your disk array software allows you to create LUNs, we recommend that you create three LUNs for the membership partitions. Each LUN should be a minimum of 8 MB in size. If you are unable to create LUNs on your disk array, you can use regular disk partitions for the membership partitions.

You *must* create partitions (using fdisk or a similar tool) on the LUNs or disks that will be used as membership partitions. It is important to partition the disks or LUNs appropriately before the membership partitions are installed. If you later need to repartition a disk containing a membership partition, you will need to stop HP Clustered File System before you change the layout. While the cluster is stopped, you will not be able to access other disks in the cluster.

5. Configure Node Settings

At this point, you are logged into a SuSE Linux Enterprise Server 9 (SLES 9) with the KDE desktop environment. If you are familiar with SLES 9, you can use any of the system configuration tools at your disposal.

To simplify the task of configuring the system, use the HP StorageWorks Clustered File System Control Panel. The GUI is a simple sequenced launch pad that guides you through the required system configuration steps. You can also use the GUI tool later to change configuration settings.

To start the tool, single click on the HP Clustered File System Control Panel icon on the desktop. Follow the steps outlined by the tool. By clicking the links in the tool, standard SLES 9 management tools are launched to perform the required configuration.

A non-graphical command line version of this configuration tool is also available. To use it from a shell window, enter the command:

```
/opt/hpcfsmgmt/bin/nasconfig
```

The following list contains the steps that the command line and graphical tools guide you through:

- a. **Set the date and timezone for the system.**
- b. **Configure Networking.** Use the information from the pre-planning worksheet to configure the three network connections on the server. Be sure to correctly configure the public and private network as discussed in Chapter 2.

NOTE: The IP address for these cards is not used by the clients. Rather, virtual IP addresses are used and assigned later during the configuration of the HP Clustered File System software.

NOTE: You must set an IP address for an Internet gateway even if you are not using one. In that instance, you can set the gateway IP to point to any of the NICs on the server.

CAUTION: Do not set up these two network connections with any kind of bonding driver as such configurations are currently unsupported.

- c. **Configure Host Names.** Provide the host names for all the network connections. It is recommended that local /etc/host names be used to speed any name resolutions. The Linux nsswitch capability has been pre-configured to look in /etc/hosts before doing any DNS name resolution. It may be simpler to create an /etc/hosts file in a text editor and then copy it to all the nodes in the cluster.

- d. **Configure Network Time Protocol (NTP).** HP recommends that you use an NTP server to ensure that the files on the cluster have consistent timestamps. Use a machine outside of the HP StorageWorks EFS Clustered Gateway.
- e. **Set up HP Management Home Page (Insight Manager) home page.** The HP StorageWorks DL380-SL Clustered Gateway comes with pre-installed system management tools. These tools can be used standalone on each server. However, they are designed to function in the larger HP-SIM and/or HP StorageEssentials management framework. For more information on HP-SIM management, go to:

[http://h18013.www1.hp.com/products/servers/management/hpsim/index.html.](http://h18013.www1.hp.com/products/servers/management/hpsim/index.html)

HP strongly recommends that you enable these Insight Manager agents by clicking the Insight Manager link. This allows you to bring up the System Management home page. A new window is displayed and asks a number of configuration questions. HP recommends that you answer yes to enable all agents. Information such as a SNMP trap destination, community strings etc will be required from the HP Pre-Planning Worksheet.

Test both local and remote access to the Management Home Page. A link exists on the desktop as well as in the Clustered File System Control Panel. The URL to use is for a browser from another system:

`https://<hostname>:2381`

- f. **Configure iLO card.**

The iLO card should have been configured during the initial hardware setup and server startup process. If this step was omitted, perform it now.

- Start by pointing a browser at the HP System Management Home Page (see the step immediately preceding this one):
- There is also a link labeled 'Lights-Out' on the left side of the Insight Manager home page. Click the link to launch the management page.

- Configure the iLO with the IP address, and add the new user and password recorded on the <pre-planning worksheet>

g. Repeat the above tasks for all the servers in the cluster.

6. Verify the SAN Configuration

This step verifies that the SAN devices are configured appropriately and can be viewed from the servers that will be in the cluster. You will need to perform this step on each server.

To begin, load the drivers provided with HP Clustered File System:

```
# /etc/init.d/pmxs load
```

Next, run the HP Clustered File System **sandiskinfo** utility, which probes the devices. It also provides a list of the devices, including configuration information.

```
# /opt/hpcfs/bin/sandiskinfo -ual
```

Review the list generated by **sandiskinfo** to verify that the SAN is configured the way that you expect.

7. Run the mxcheck Utility

This utility should be run on each server. It verifies that the server's configuration meets the requirements for running HP Clustered File System. Issue the following command:

```
# /opt/hpcfs/bin/mxcheck -l -p -t - install
```

The utility performs checks such as the following: operating system version and configuration; available physical memory and disk space; HBA driver versions; FibreChannel switch versions; network addresses. The utility also attempts to access gateways and FC switches.

Output from the utility appears on the screen and is also written to the */var/hpcfs/mxcheck* directory. We recommend that you fix any problems identified by **mxcheck** before you configure HP Clustered File System.

8. Run the **mxconfig** Utility

The section is presented assuming the use of the **mxconfig** utility. You can also perform all of these steps using the graphical Management Console (**mxconle**). If you use the **mxconsole** tool, select **Configure** to start the configuration process. The default login information for **mxconsole** is “admin” for the username, and “hpinvent” for the password. Refer to Appendix A on page 54 for details on using the graphical Management Console for doing this part of the configuration.

The **mxconfig** utility performs some initial configuration of HP Clustered File System. It allows you to upgrade the license file, to select a cluster password and a Network Authentication Secret password, to select a fencing method, to specify the FibreChannel switches connected to the nodes in the cluster, and to select the LUNs or disk partitions to be used as HP Clustered File System membership partitions.

You can run **mxconfig** on one server and then export the resulting configuration to the other servers. You must be *root* to run this utility. (To see help information, type **mxconfig --help**.)

```
# /opt/hpcfs/bin/mxconfig
```

You can also click on the link in the HP Clustered Gateway Control Panel titled “Initial HP Clustered File System [mxconfig].”

When you invoke **mxconfig**, the following Welcome window appears.

NOTE: If **mxconfig** does not display properly, increase the size of your window and then run **mxconfig** again. Setting the TERM variable to **xterm** may also make the display easier to read.

Welcome to mxconfig

Welcome to mxconfig. You may abort mxconfig at any time by pressing the <esc> key.

< OK >

On windows that require input, use the Tab key to move between OK and Cancel or between Yes and No. Press the Enter key to go to the next window. Press the Escape key to abort the **mxconfig** utility.

HP Clustered File System License File

The HP Clustered File System product ships with the required license file pre-installed. Complete the following step to upgrade a demo license file to a permanent license or to upgrade a license file to enable future new product capabilities.

HP Clustered File System can be used with either a temporary or a permanent license. In this case, the license is provided in a separate license file. (The file must be present on the server you are using to run **mxconfig**.) To install the license, select Yes on the Replace License File window. When prompted, enter the location of your license file. **mxconfig** will copy this file to */etc/hpcfs/licenses/license*.

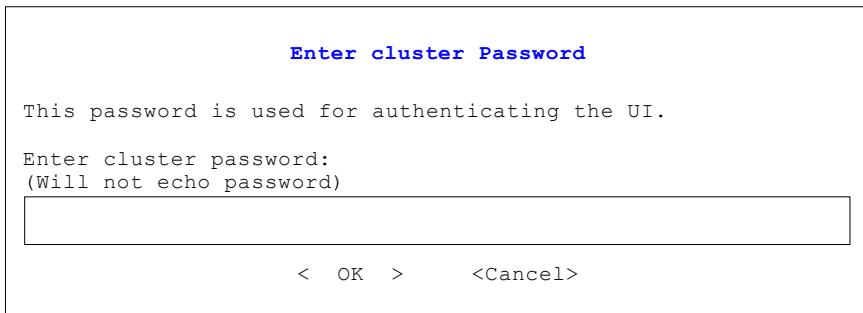
New License File Location

Enter Location of New License File:

< OK > <Cancel>

Specify the Cluster Password

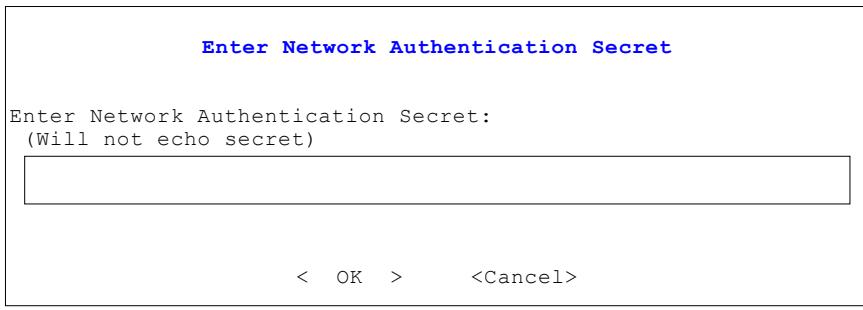
The password you specify on the Enter cluster password window is assigned to user *admin*. (Log into the Management Console as user *admin* to configure the cluster.) The password does not display on the window as you type it.



When prompted, re-enter the password.

Specify the Network Authentication Secret Password

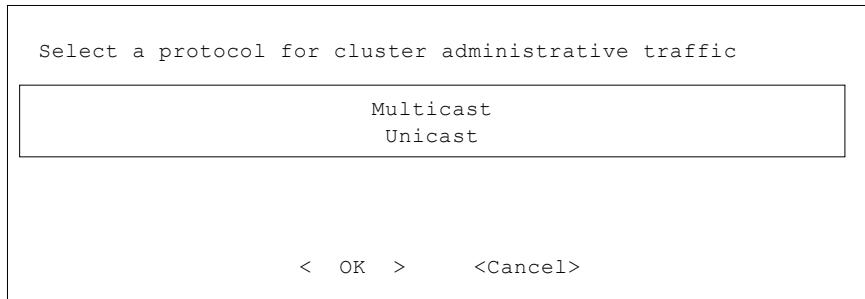
The Network Authentication Secret password provides additional security and is used by the network connections. Enter anything you want for this password. A unique password must be used when multiple clusters are being deployed to prevent cross-cluster issues.



Select the Cluster Administrative Traffic Protocol

Specify either multicast or unicast mode. Multicast mode is recommended; however, if your network configuration does not allow

multicast traffic through the network, use unicast mode. Unless you are already using multicasting in your network, HP recommends you select **Unicast** mode.



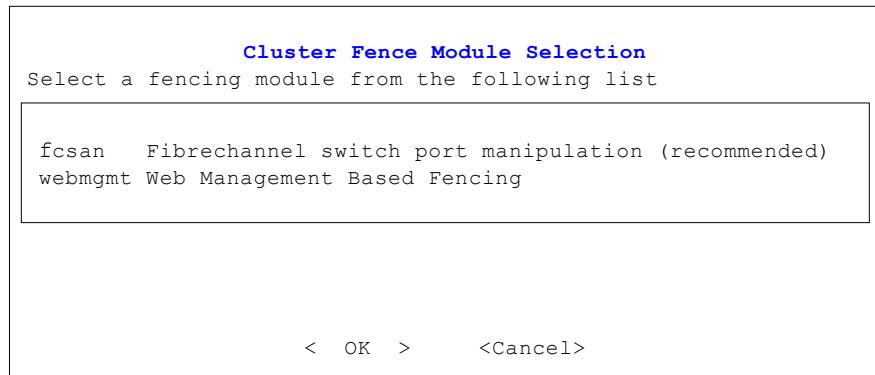
Select a Fencing Method

NOTE: The HP Clustered Gateway product has been tested with the Web Management-Based Fenced (also called server fencing). Select this option when you configure this part of the solution.

When certain problems occur on a server (for example, hardware problems or the loss of cluster network communications), and the server ceases to effectively coordinate and communicate with other servers in the cluster, HP Clustered File System must remove the server's access to filesystems to preserve data integrity. This step is called *fencing*.

There are two fencing methods:

- **FibreChannel switch port manipulation.** When a server needs to be fenced, HP Clustered File System will disable the server's access in the FibreChannel fabric. The server must be rebooted to regain access to the SAN.
- **Web Management-Based Fencing.** HP Clustered File System uses the ProLiant iLO card on the server to remove its access to PSFS filesystems. (See the *HP Clustered File System Hardware and Software Compatibility Guide* on the HP web site for a list of supported servers and firmware.)



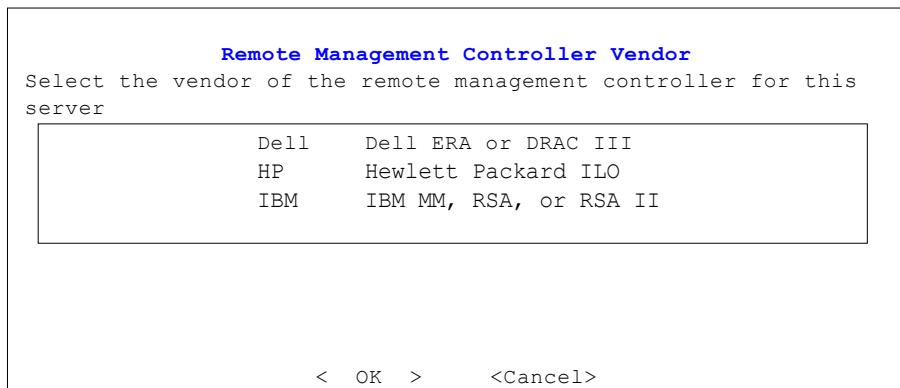
If you selected FibreChannel switch port manipulation, next go to “Configure FibreChannel Switches” on page 31.

If you selected Web Management-Based Fencing, you are prompted to supply the information described below for the iLO Card associated with the server that you are using to run **mxconfig**.

Later in this procedure the complete cluster configuration is exported to the other servers that are in the cluster. During the export, HP Clustered File System uses, where possible, the information that you specify here to configure Web Management-Based Fencing on the other servers. If a particular entry did not apply to all servers, HP Clustered File System prompts you for that information for each of the other servers.

NOTE: After HP Clustered File System is configured, you can verify that it has the correct information to fence each server. See “Verify the Fencing Configuration” on page 43.

Remote Management Controller Vendor. Select HP as the vendor for the Remote Management Controller on the server on which you are running **mxconfig**.

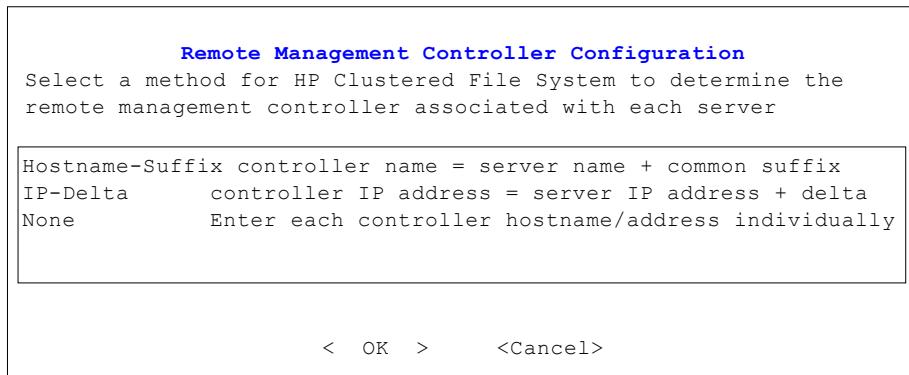


You will next be asked whether all servers in the cluster are from the same vendor.

Remote Management Controller ID. On the Remote Management Controller Configuration window, specify how HP Clustered File System should identify the Remote Management Controller associated with each server. Use one of the following methods.

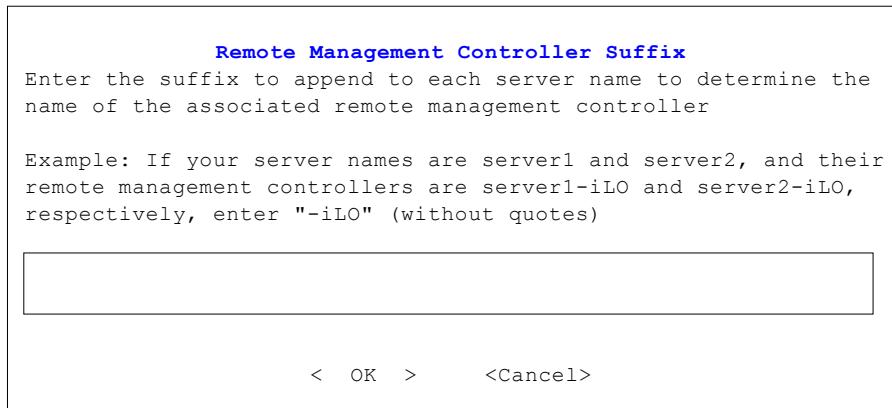
- Specify the common naming scheme that you are using for the Remote Management Controllers (either a hostname suffix or an IP address delta).

- Enter the hostname or IP address for the Remote Management Controller associated with this cluster server.



Enter the configuration information in accordance with the method that you selected.

Hostname Suffix. Specify the common suffix to append to each server name to determine the associated Remote Management Controller name. For example, if your server names are server1 and server2 and their Remote Management Controllers are server1-iLO and server2-iLO, enter **-iLO** as the suffix.



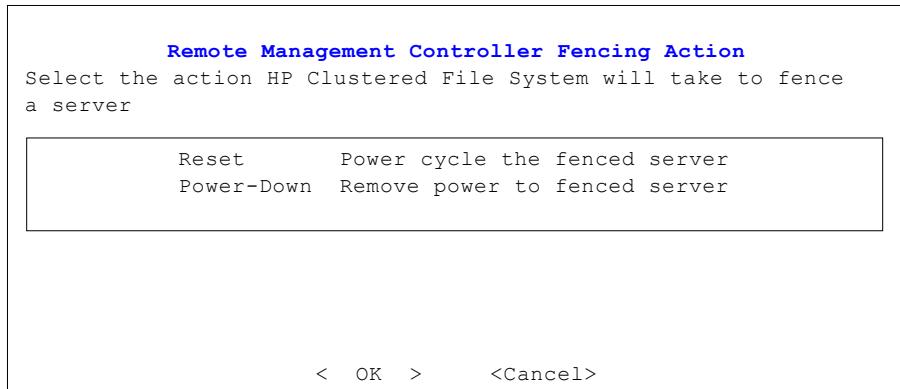
IP Delta. Specify the delta to add to each server's IP address to determine the IP addresses of the associated Remote Management Controllers. For example, if your servers are 1.255.200.12 and 1.255.200.15 and their Remote Management Controllers are 1.255.201.112 and 1.255.201.115, enter **0.0.1.100** as the delta.

<p>Remote Management Controller IP Delta Enter the delta to add to each server's primary IP address to determine the associated remote management controller IP address</p> <p>Example: If your servers are 1.500.200.12 and 1.500.200.15, and their remote management controllers are 1.500.201.112 and 1.500.201.115, respectively, enter "0.0.1.100" (without quotes)</p> <input type="text"/> <p>< OK > <Cancel></p>

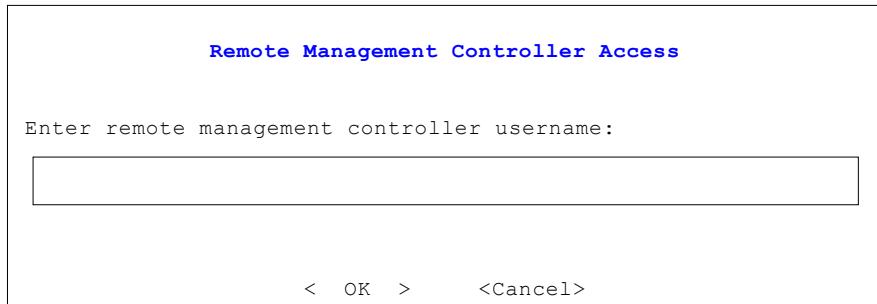
Hostname/IP Address. Specify either the hostname or IP address for the Remote Management Controller associated with the server to which you are connected.

<p>Remote Management Controller Configuration Enter the hostname or IP address of the remote management controller associated with this server.</p> <p>You will be required to enter this information for other servers in the cluster when the configuration is exported or imported to them.</p> <input type="text"/> <p>< OK > <Cancel></p>

Fencing Action. When a server needs to be restricted from the SAN, HP Clustered File System can either power-cycle the server or shut it down. Specify the method that you want to use on the following window.



Remote Management Controller Access. HP Clustered File System needs to log into a user account on the Remote Management Controller (iLO) to fence the server. The account must currently exist on the Remote Management Controller. Specify the user name for the account that you want HP Clustered File System to use.



Now specify the password associated with the user account.

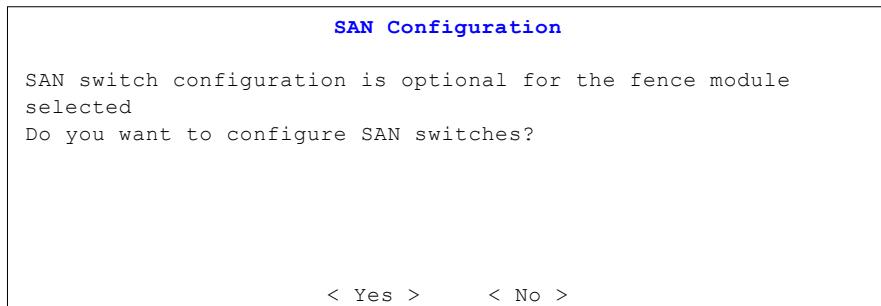
Remote Management Controller Access		
Enter remote management controller password: (Will not echo password)		
<input type="text"/>		
< OK > <Cancel>		

You are now asked whether the same username and password are used by all of the Remote Management Controllers in the cluster.

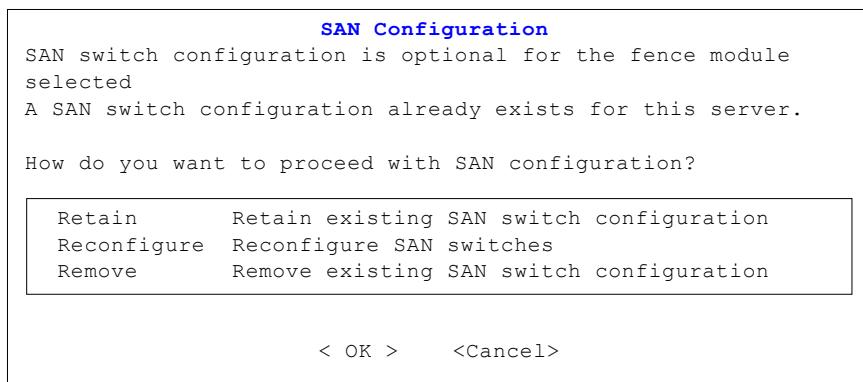
Remote Management Controller Access		
Do all nodes of this cluster share the same remote management controller username and password?		
< Yes > < No >		

SAN Configuration. You are prompted whether you want to configure the SAN switches in HP Clustered File System. This step is optional for Web Management-Based Fencing configurations; however, if the switches are configured the Management Console can display the switch ports used by the SAN.

If the SAN switches have not previously been configured in HP Clustered File System, you will see the following window.

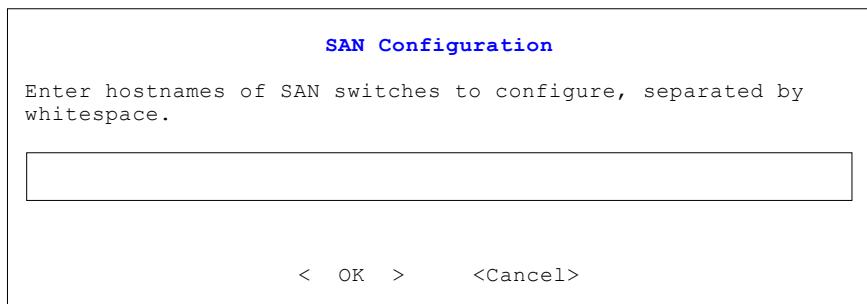


If HP Clustered File System detects an existing SAN switch configuration, the following window will appear. You can retain, remove, or modify the existing configuration.



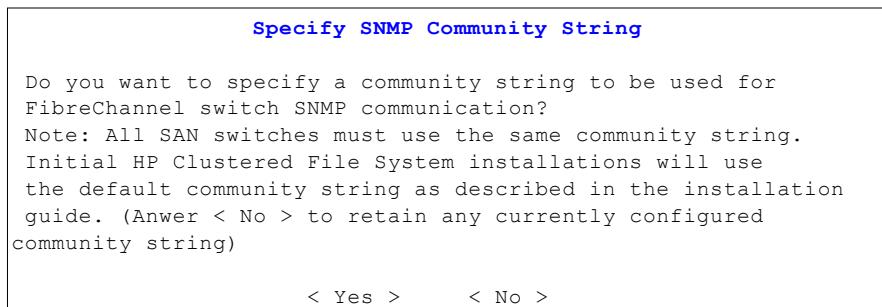
Configure FibreChannel Switches

On the SAN Configuration window, specify the hostnames of the FC switches that are directly connected to the nodes in the cluster. (If you are using Web Management-Based Fencing, you will see this window only if you chose to configure or reconfigure SAN switches.) Even when Web Management-Based Fencing is used, configuring the access to the switch can be helpful for troubleshooting. HP recommends that you provide this information.

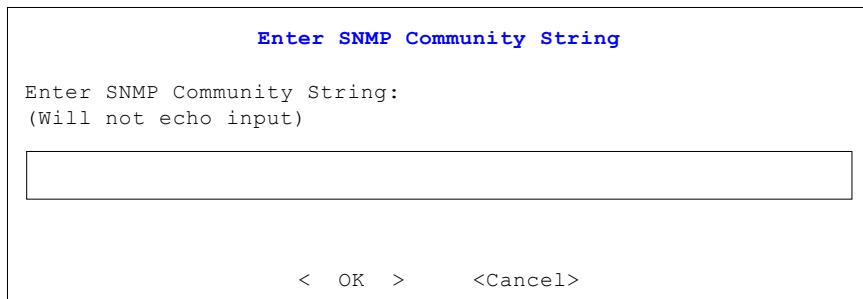


Set the SNMP Community String

The SNMP community string must be set to the same value on HP Clustered File System and on the SAN switches managed by the cluster. The default SNMP community string for HP Clustered File System is **private**. If you want to use a custom community string instead, select **Yes** on the Specify SNMP Community String window.



If you selected **Yes** on the Specify SNMP Community String window, type the appropriate string on the Enter SNMP Community String window that appears next.



Modify Disk Partitioning for Membership Partitions

HP Clustered File System uses a set of membership partitions to control access to the SAN and to store the device naming database, which includes the global device names that HP Clustered File System assigns to the SAN disks placed under its control.

Select the partitions that should be used as membership partitions. HP Clustered File System can use either one or three membership partitions. **To ensure that a membership partition is always available, HP strongly recommends that you select three partitions.**

A membership partition can be located either on a LUN or on a regular disk partition.

- LUNs. Create a partition on each LUN. Answer yes on the Create Membership Partitions window and then use **fdisk** to create a partition on each LUN. These partitions can then be used for the membership partitions.
- Regular disk partitions. To modify the partitioning on the disks that will contain the membership partitions, answer Yes on the Create Cluster Membership Partitions window to invoke the **fdisk** utility on those disks. (8 MB is adequate for a membership partition.)

Ensure that the partition layout is adequate for your future needs. When you start HP Clustered File System, the disks containing the membership partitions will be imported into the cluster. These disks cannot be removed from the cluster while HP Clustered File System is running. If you need to repartition a disk containing a membership partition, you will need to stop HP Clustered File System before you change the layout. While the cluster is stopped, you cannot access other disks in the cluster.

Create Cluster Membership Partitions

Do you want to use fdisk to create new partitions for cluster membership?

< Yes > < No >

If you chose to run **fdisk**, you will be asked to select the disk or LUN where you want to modify the partitioning. The **fdisk** utility will then be invoked on that disk or LUN. When you have completed your changes, you can select another disk or LUN to be modified with **fdisk**. When you have finished modifying partitions, select Cancel.

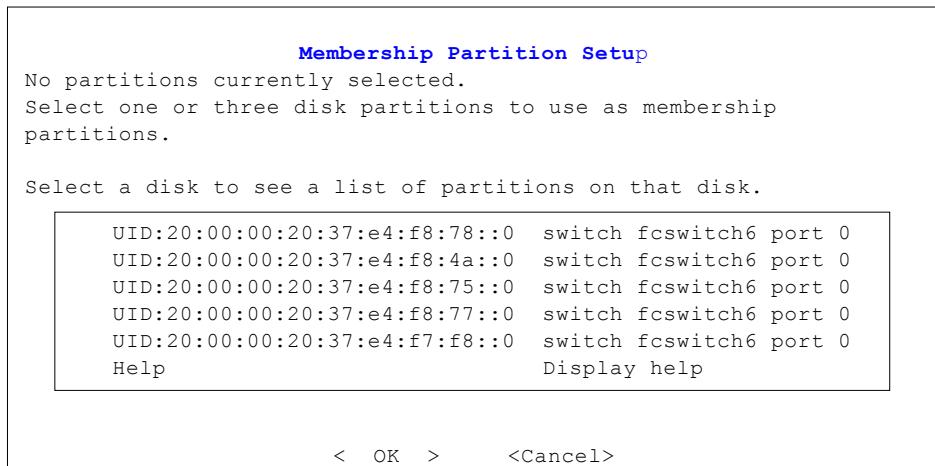
NOTE: When you use **fdisk**, the modified partition table is visible only on the server where you made the changes. When you start HP Clustered File System, disks or LUNs with membership partitions are imported into the cluster automatically. The revised partition table will then be visible to all of the servers.

NOTE: When using the graphical **mxconsole** tool, you cannot run **fdisk** directly from the GUI. Rather, you must start it manually from a shell window.

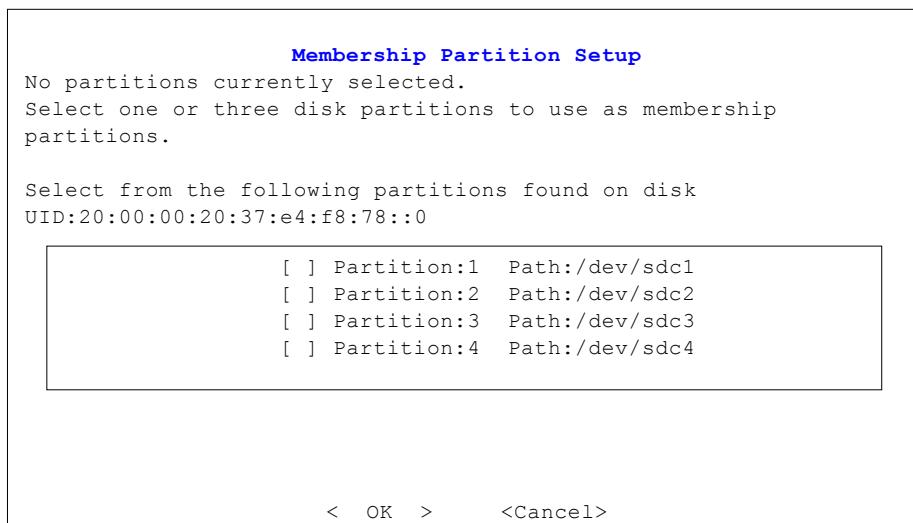
Create Membership Partitions

The Membership Partition Setup window prompts you to select a disk where you want to create a membership partition. Use the arrow keys to move to the appropriate disk or LUN and then select OK.

When a disk or LUN is highlighted, the vendor, model, and capacity of the disk or LUN will be displayed at the bottom of the window. The local path to the disk or LUN also appears.



After you select OK, the Membership Partition Setup window shows the partitions on the disk or LUN. Use the arrow keys to move among the partitions. When a partition is highlighted, its size is displayed at the bottom of the window. To select a partition to be used as a membership partition, move to that partition and then press the spacebar. One MB is adequate for a membership partition.



The partition you selected is displayed on the Membership Partition Setup window. Choose OK to complete the selection. Repeat this procedure to select two more membership partitions. HP recommends that the partitions be on different disks. You can deselect and reselect partitions as necessary.

NOTE: If you select a large partition, you will see a warning saying that the partition is much larger than needed. You can either ignore the warning or change your selection to a smaller partition.

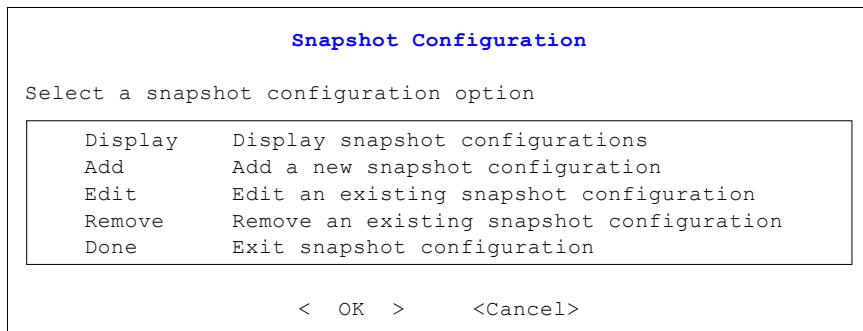
To complete the operation, select “Done,” which appears after the list of disks on the Membership Partition Setup window. (You may need to scroll down to reach “Done.”)

mxconfig next asks you to confirm your selections and then initializes the membership partitions.

Add a Snapshot Configuration

HP Clustered File System supports taking hardware snapshots of PSFS filesystems. (The filesystems must be located on storage arrays supported for snapshots.) If your HP Clustered File System license includes the snapshot feature, the Snapshot Configuration menu is displayed next.

This menu provides options to display existing snapshot configurations, to add a new snapshot configuration, and to edit or remove existing snapshot configurations.



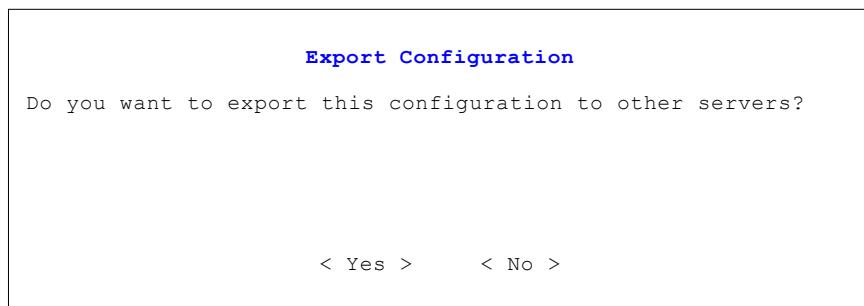
To add a new snapshot configuration, select Add. Then specify the requested information on the screens that appear next.

For the HP EVA Management Appliance, you are prompted for the hostname or IP address of the appliance associated with the cluster, and the username and password to use when accessing the appliance.

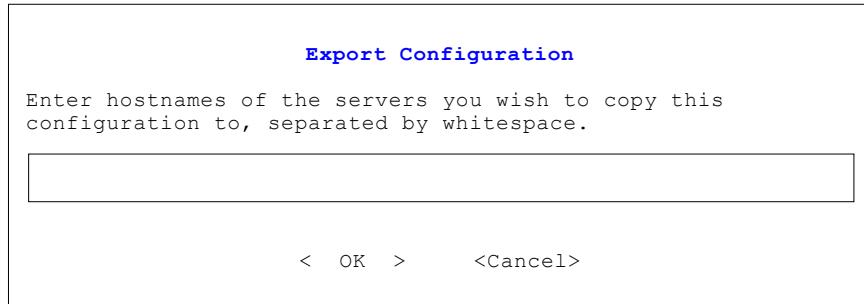
NOTE: The username and password are for the CommandView EVA application that runs in the management appliance. This is *not* the same username and password as for the overall management appliance! The default username and password are both "administrator." Ask your administrator for the current settings.

Export the Configuration

The initial configuration is now complete. **mxconfig** can export this configuration to the other servers in the cluster.



If you chose to export the configuration, type the names of the servers that you want to receive the configuration on the Export Configuration window. Use white space to separate the names.



mxconfig uses **ssh** as user *root* to copy the configuration to each server. Depending on your ssh configuration, you may be prompted for the root password one time on each host (this is the default behavior).

If you are using Web Management-Based fencing, you may be asked for additional information about each server. Based on your fencing configuration, the information can include any or all of the following: the vendor/type for the Remote Management Controller, the IP address/hostname of the Remote Management Controller, the username and password for the Remote Management Controller.

9. Start HP Clustered File System on One Server

Start HP Clustered File System on one server that will be in the cluster.

Do not start HP Clustered File System on the other servers.

```
# /etc/init.d/pmxs start
```

NOTE: You can also use the shortcut on the HP Clustered Gateway Control Panel to perform this task.

NOTE: If you are using the Management Console (**mxconsole**) GUI to perform this part of the configuration, select the Cluster Wide Configuration tab. From that screen you can add servers, start servers, and so on.

This script calls the **mxinit** utility, which starts the HP Clustered File System cluster processes in the correct order and loads the kernel modules.

10. Start the Management Console

Start the Management Console on the server on which you started HP Clustered File System. First start the X windowing environment and then type the following command.

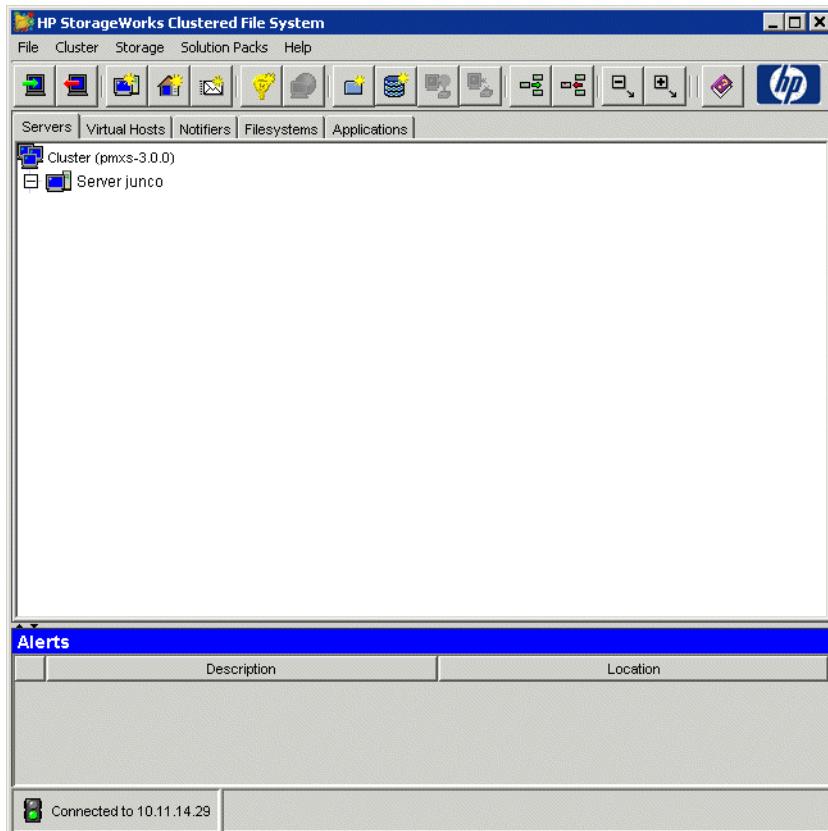
```
$ mxconsole
```

You can also launch the Management Console by clicking the corresponding desktop icon, or on the task within the HP Clustered Gateway Control Panel.

When the Login window is displayed, enter the name or IP address of the server, leave *admin* as the User and enter the password that you selected earlier. Click the Login button.



The Management Console window is displayed. The server where you started HP Clustered File System is currently the only server in the cluster.



11. Add the Remaining Servers to the Cluster

To add another server to the cluster, select Cluster > Server > New and enter the name or IP address of the server on the New Server window.



Repeat this procedure to add the remaining servers to the cluster. The Management Console indicates that the servers are down because HP Clustered File System is not yet running on them.

12. Start HP Clustered File System on the Remaining Servers

You can start the required services in one of two ways. The easiest way is to use the Management Console (**mxconsole**) in its **Configuration** mode. Select the Cluster Wide Configuration tab, and look for **Start Service** button to start services on other nodes without having to log into each node. If you choose to log into each node, run the following script on each server that you added to the cluster.

```
# /etc/init.d/pmxs start
```

NOTE: When HP Clustered File System is started on a new server, ClusterPulse updates the SAN information on that server. During the update, the Management Console may be unable to report SAN information. You can ignore this; the cluster is operating correctly and the SAN information will reappear within a few seconds.

The installation is now complete. See the *HP StorageWorks Clustered File System Administration Guide* for information about configuring the cluster.

13. Configuring HP CFS for Public versus Private Network

After a cluster has been formed, configure it to correctly use the private network for intra-cluster network communication. The HP CFS software does not identify which network connection is the “private” network. You must indicate which network is considered to be the private one. Failure to perform this step could allow the public network to be selected which could adversely affect performance. It is important to correctly complete this task. To do this, follow these instructions:

1. Use the HP Clustered File System Console to:
 - a. Select the 'Servers' Tab

- b. Repeat the following step for each public network:
- c. Right click on the network interface corresponding to the public network
- d. Select 'Discourage Admin Traffic'

2. Use the command line to discourage intra-cluster traffic on your the public network connections (there are typically two of them):
 - - mx netif noadmin eth1
 - - mx netif noadmin eth2

NOTE: This assumes that eth1 and eth2 correspond to your public network. To verify this, you can use 'ifconfig ethN' to look at network settings.

Discouraging traffic on the public interfaces gives preference to the private network. However, if the private network fails, the public networks are used as an alternative means of communication. For more information, see the *HP StorageWorks Clustered File System Administration Guide* topic "Allow or Discourage Administrative Traffic."

14. Set Up NFS file services

The NFS file services setup is described in the *HP StorageWorks Clustered File System for Linux: NFS Configuration*. This software provides the highly scalable, highly available NFS file serving capabilities.

A quick link to the HP Clustered File System Software management tool is provided on the desktop. Click the icon to start the configuration tool.

15. Save System State

It is possible to create a system state backup. This is useful if the server ever has to be rebuilt. The system save state will contain all the configuration information required to return the server to its current state. HP strongly recommends that you do this to all nodes on the system.

However, it is probably best that the HP EFS Clustered Software be configured first. After the cluster is fully functional, perform a system state backup.

The save state data set is fairly small (around 10 Mb). Save the generated data set someplace other than the disc of the system being backed up. Whenever base system configurations change, generate a new system save state.

The configuration information for the HP EFS Clustered Software is kept on management partitions on the SAN storage. Changes to this configuration data do not require a new system save state to be generated.

The process of creating the save state dataset is initiated from the graphical management tool.

4

Post Setup Procedures

This chapter discusses tasks to perform after you complete the initial setup. The following procedures are discussed.

- Verify the Fencing Configuration
- Install the Management Console on Other Hosts
- Disaster Recovery
- CIFS File Serving
- Quota Management

Verify the Fencing Configuration

The **mxfence** command-line utility can be used to verify that HP Clustered File System has the information needed to fence a server. This utility is intended for Web Management-Based fencing configurations where you must specify the hostname or IP address of the Remote Management Controller associated with the server. You can use **mxfence** to verify that HP Clustered File System has the correct information. The server must be up when you use the utility.

/opt/hpcfs/sbin/mxfence <server hostname or IP address>

When you run **mxfence**, HP Clustered File System uses the specified hostname or IP address to access the Remote Management Controller. The server is then either power-cycled or shut down in accordance with the method you selected when you configured the fencing module.

If the correct server is not fenced, the hostname or IP address specified for the Remote Management Controller is probably incorrect. Run **mxconfig** on the server with the incorrect information and then modify the information as necessary.

Install the Management Console on Other Hosts

For best performance, the cluster should be monitored from a separate administrative station rather than from a server in the cluster. The Management Console software can be installed on either Linux or Microsoft Windows systems.

Management Console for Linux Systems

System Requirements

Systems used for the Management Console must meet the following requirements:

- Pentium-compatible processor or later.
- 64 MB of memory.
- SuSE Linux Enterprise Server Version 7, 8, or 9, Red Hat Linux 7.2 or 7.3 (Server or Workstation installation), RHAS 2.1, RH AS/ES 3.0.
- On Red Hat systems, the “compat-libstdc++” package must be installed.
- A windowing environment such as the X Window System must be installed and configured.

Install and Start the Management Console

To install the Management Console on a Linux system, locate a copy of the installable RPM package. There are two choices:

- Mount the HP Clustered File System Documentation CD and look in the **mxconsole** subdirectory.

- Obtain a copy of the installable RPM package from the /opt/hpcfsmgmt/tools directory on a running cluster node.

After you have a copy of the installable package and made it available on the system where you wish to install the Management Console, run the following command:

```
# rpm -i <location of the package>/mxconsole-3.0.0-<xxxx>.i386.rpm
```

To start the Management Console, type the following command:

```
$ mxconsole
```

When the Connect to Cluster window appears, you can log into the cluster.

Management Console for Windows Systems

System Requirements

Systems used for the Management Console must meet the following requirements:

- Pentium-compatible processor or later
- 64 MB of memory
- Windows 2000, Windows XP, or Windows 2003

Install and Start the Management Console

To install the Management Console on a Window system, locate a copy of the installable MSI file. There are two choices:

- Mount the HP Clustered File System Documentation CD and look in the **mxconsole** subdirectory.
- Obtain a copy of the installable MSI file from the /opt/hpcfsmgmt/tools directory on a running cluster node.

After you have obtained a copy of the installable package and made it available on the system where you wish to install the Management Console, double-click the file:

```
<location of file>/mxconsole/mxconsole-3.0.0-<xxxx>.msi.
```

To start the Management Console, select Start > Programs > HP Management Console > Management Console.

Disaster Recovery

HP StorageWorks DL380-SL Clustered Gateway servers in the cluster are highly reliable. The hardware has redundant components that should keep a server running in most scenarios. If you must replace a server or re-image the operating system and software, a quick restore DVD is included with the product for this purpose.

CAUTION: The quick restore DVD restores a cluster node to its original factory state. This is a destructive process that completely erases all of the data on the two system discs.

Re-Image the Cluster Node

To re-image a node, complete the following steps:

1. Install the quick restore DVD in the DL380's DVD ROM drive.
2. Reboot the system.
3. Type "R" (for restore) when prompted to do so or "E" to exit.

NOTE: This completely erases all of the data on both mirrored hard drives.

When the restore process is finished, the DVD ejects and the system automatically reboots. You can now perform a node state restore, if you performed a system state backup previously.

Restore the Node State

If you performed a system state backup and saved the backup to a safe location (refer to the step "15. Save System State" on page 41), you can use the backup to restore the node to the state represented by the backup. This includes restoring items such as network settings and any other customizations done on the system.

NOTE: Performing a system state restore does not cause the node to re-join its previous cluster. Use the HP Clustered File System Console to add the node back into the cluster.

1. After the system reboots, log in with the username `root` and password `hpinvent`.
2. Use the "Recover State" link the HP Clustered Gateway Control Panel to restore a previously saved system state.
3. Add the server back into the HP StorageWorks EFS Clustered Gateway.

CIFS File Serving

The HP StorageWorks EFS Clustered Gateway is for use in NFS environments. However, the solution supports mixed protocol file serving as well. In particular, the CIFS protocol is supported by the Samba package inherent in most Linux distributions including SLES 9.

Use the standard Samba tools as well as the and SLES “YaST2” tool to configure Samba shares.

Samba does not support clustered environment. The following limitations therefore apply to its use in the HP Clustered File System:

- Samba may be run on one or more nodes of a cluster.
- If Samba is run on multiple nodes, each Samba instance must export a *unique* part of the file system. Failure to follow this guideline could result in corrupted data (no cross node locking is possible with Samba).
- If the Samba service is to be highly available (see below), one or more cluster nodes must be reserved as Samba failover nodes.

It is possible to create a highly available Samba configuration.

Specifically, if the server that hosts the Samba service fails, it can be automatically started up on another node in the cluster. This functionality is provided by a set of *unsupported* scripts that work in conjunction with the HP EFS Clustered Software. These scripts and instructions for using them are located in the

`/opt/hpcfs/contrib/SAMBA/directory`.

Quota Management

The HP Clustered File System supports soft user quotas through the use of a set of scripts. These scripts are unsupported by HP but can be used effectively used to monitor storage consumption. They are found in `/opt/hpcfs/tools`.

The following list is an overview of the included scripts, and what they are used for:

- `dubyuid` — An executable that takes one or more subtrees as arguments and prints out the disk usage of each user in 512-byte blocks.
- `updatesqdata` — A simple script to be run from cron(1M) that called `dubyuid` for each tree listed in `/etc/hpcfs/softquotas/fslist`. Output is placed in `/var/hpcfs/softquotas/quotadata`.
- `runquotas` — A simple script that checks the `quotadata` output against a list of user quotas in `/etc/hpcfs/softquotas/quotas` and prints out a line for each user who is over quota.
- The `fslist` file is a simple text file. Each line should contain a path to a top-level directory under which disk usage is to be measured. The `quotas` file is a simple text file. Each line should contain two fields delimited by white space. The first field contains the username, and the second field contains the quota for that user expressed in 512-byte blocks.

Additional information can be found in the file `/opt/hpcfs/tools/softquota.readme`.

5

Other Procedures

This chapter describes how to perform the following procedures:

- Add a new server to a cluster
- Upgrade the license file
- Change the Fencing Method
- Uninstall the Management Console
- Set HP Clustered File System Parameter for Falconstor

Add a New Server to a Cluster

Use the following procedure to add a new server to an existing cluster.

1. **Set up HP Clustered File System.** Perform the procedure in Chapter 2 up to the step to run **mxconfig**.
2. Use the following command to invoke the **mxconfig** utility:

```
# mxconfig --import
```

The **--import** option allows you to import the initial HP Clustered File System configuration defined with **mxconfig** from an existing cluster server to the new server. You will be asked to specify the hostname of the server from which the configuration should be copied.

Depending on your fencing method, you might also need to specify the hostnames of the FC switches that are directly connected to the servers in the cluster.

If you are using Web Management-Based fencing, you may be asked for additional information about the server. Based on the fencing configuration, the information can include any or all of the following: the vendor/type for the Remote Management Controller, the IP address/hostname of the Remote Management Controller, the username and password for the Remote Management Controller.

mxconfig then copies the initial configuration to the new server.

3. From the Management Console, add the new server to the cluster.
4. Start HP Clustered File System on the server:

```
# /etc/init.d/pmxs start
```

When a new server is added, ClusterPulse updates the SAN information on that server. During the update, the Management Console may be unable to report SAN information such as the status of PSFS filesystems. You can ignore this; the cluster is operating correctly and the SAN information will reappear within a few seconds.

Upgrade the License File

If you purchase an upgraded license for your cluster, you will be sent a new license file. You will need to manually copy this file to */etc/hpcfs/licenses/license* on each server.

On each server, HP Clustered File System reads the license file upon startup and at 15-minute intervals. Until the new license file is read on all servers, you may see alerts on the Management Console reporting that the license file does not match on all servers.

Change the Fencing Method

This procedure describes how to change the cluster fencing method on servers that are already running HP Clustered File System 3.0.0.

1. Stop HP Clustered File System on all servers in the cluster:

```
# /etc/init.d/pmxs stop
```

2. Clear the host registry on the cluster. On one server, issue the command **mxmpconf** at the operating system prompt. Then select “Repair” from the Main Menu. On the Repair Menu, select the option to clear the host registry.
3. Run **mxconfig** from one server in the cluster. (See installation step “7. Run the mxcheck Utility” on page 19.) Skip the opening windows and then configure the appropriate fencing method.
4. When you have completed the fencing configuration, skip the membership partition configuration windows and then export the new configuration to the other servers in the cluster.
5. Start HP Clustered File System on each server:

```
# /etc/init.d/pmxs start
```

6. Verify the fencing configuration. See the procedure “Verify the Fencing Configuration” on page 43.

Uninstall the Management Console

To uninstall the Management Console on a Linux system, type this command:

```
# rpm -e mxconsole-3.0.0-<xxxx>
```

To uninstall the Management Console on a Windows system, select Start > Settings > Control Panel > Add/Remove Programs and remove the application.

Set HP Clustered File System Parameter for FalconStor

This step is needed only if a FalconStor device is included in your SAN configuration.

The HP Clustered File System **psd_round2_delay** parameter must be set manually on each server. This parameter is located in the *scl.conf* file, which by default is located at */etc/hpcfs/conf/scl.conf*.

To set the **psd_round2_delay** parameter, locate these lines in the *scl.conf* file:

```
#psd_round2_delay
#      (This parameter is only used for Linux; it has no effect on Windows.)
#      Set the delay time (in seconds) that the psd driver will wait before
#      doing a second round of retries for an I/O that has already errored
#      out on all available paths. This used for devices that require
#      time to recover from errors (i.e.: FalconStor).
#
#      (Setting this value to -1 will disable round2 processing altogether.
#      The default value is -1, no round2 processing.)
#
#psd_round2_delay      -1
```

On the last line, remove the # sign preceding **psd_round2_delay** and replace -1 with the number of seconds to wait before the psd driver retries the I/O. The recommended value is 45 seconds.

```
psd_round2_delay      45
```

A

Configure the Cluster from the Management Console

Cluster Configuration

This appendix describes how to run the **mxconfig** utility from the Management Console. A windowing environment must be installed on the server where you are running the Management Console.

Required Information

You will be asked for the following information:

- The IP address or hostname of each server that will be in the cluster.
- The location of your HP Clustered File System license file.
- A secret network key that provides security for network communications among the cluster servers.
- An administrative password for configuring the cluster.
- The fencing method that you want to use (either FibreChannel switch-based fencing or Web Management-Based Fencing via server reset/shutdown).
- The IP address or hostname of each FibreChannel switch that is included in the SAN. (This step is required for FibreChannel switch-based fencing and optional for Web Management-Based Fencing via server reset/shutdown.)

- The partitions or LUNs that you want to use for membership partitions.

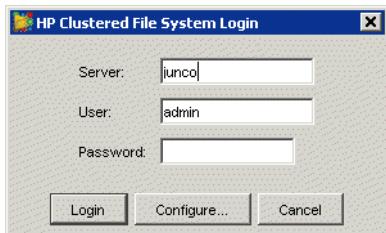
Configuration Procedure

Connect to the Management Console

To start the Management Console, first start the windowing environment on the server and then type the following command:

```
$ mxconsole
```

On the Login window that appears next, enter the IP address of the server, type **admin** for the user and **hpinvent** the password, and then click the Configure button.

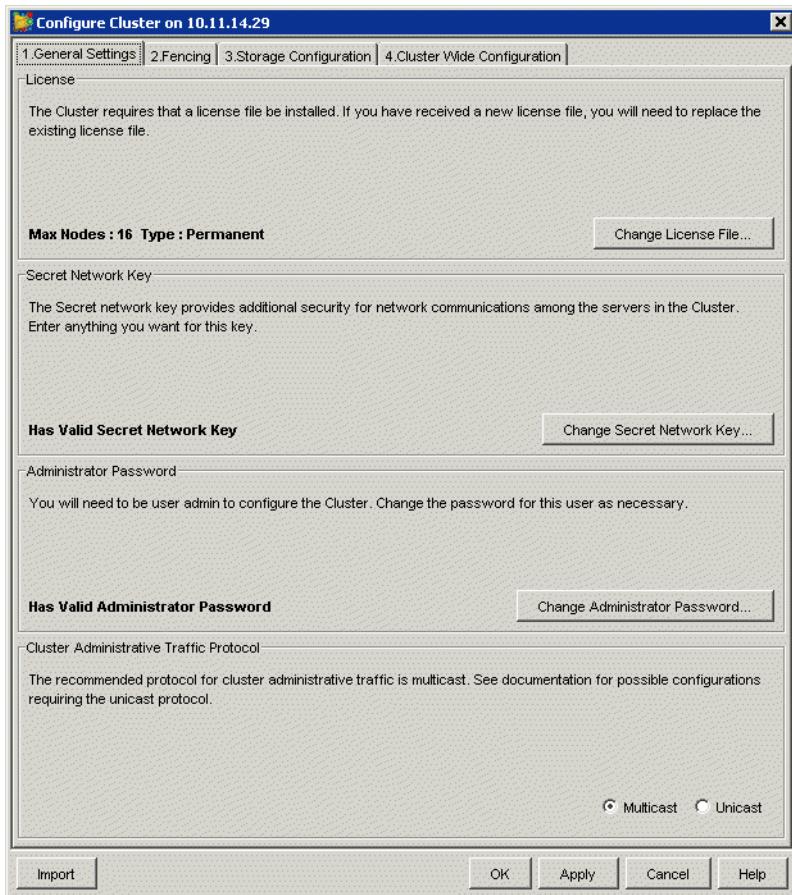


NOTE: If you click the Login button, an error message appears and you will be asked whether you want to set up the cluster. Click Yes to configure the cluster.

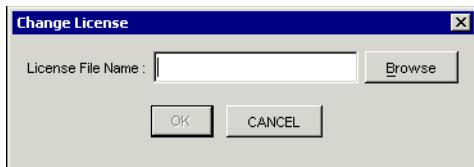
The Configure Cluster window then appears. You will need to specify information on the tabs in this order: General Settings, Fencing, Storage Configuration, Cluster Wide Configuration.

General Settings

This tab asks for general information needed for cluster operations.



1. **License.** HP Clustered File System can be used with either a temporary or a permanent license. The license is provided in a separate license file. (This file must be present on the server that you are using to connect to the Management Console.) To install the license, click the Change License File button. Type the location of the license file or use the Browse button to locate it. The license file will be installed at `/etc/hpcfs/licenses/license`.



2. **Secret Network Key.** This password is required. It provides additional security for network communications among the cluster servers. To set this key, click the Set Secret Network Key button. You can enter anything you want for this password.



3. **Administrator Password.** You will need to be HP Clustered File System user *admin* to configure the cluster. By default, the password for this user is set to *admin*. If you want to change the password, click the Change Administrator Password button.

The password can be up to 32 characters long with no white space. The allowable characters are ASCII 33–122.



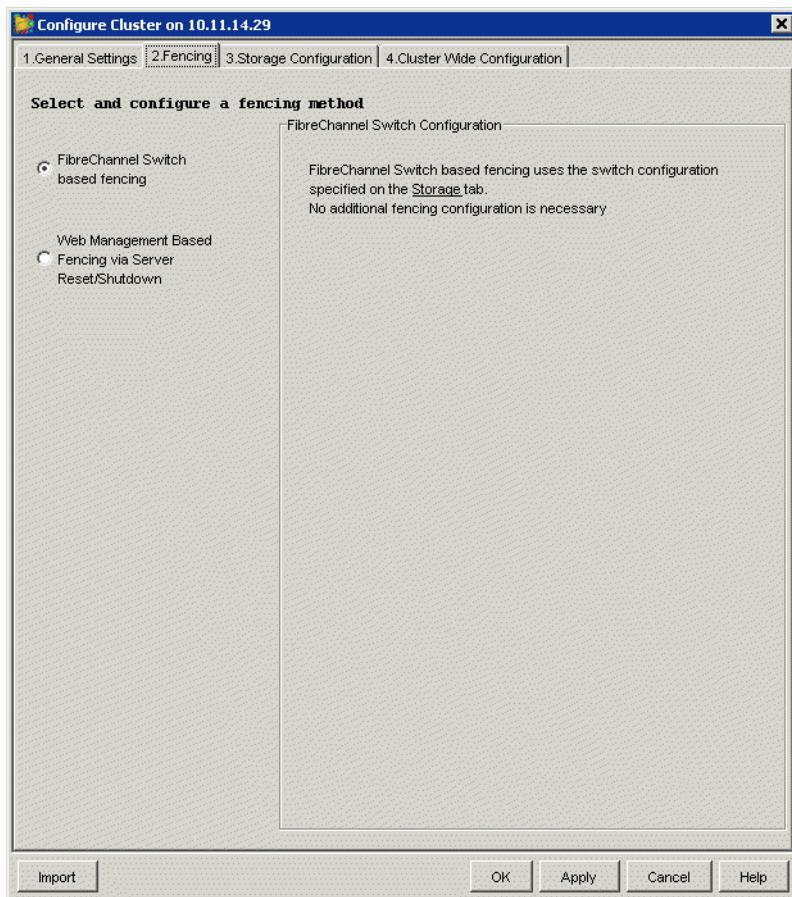
4. **Cluster Administrative Traffic Protocol.** Select either multicast or unicast mode. Multicast mode is recommended; however, if your network configuration does not allow multicast traffic through the network, you will need to use unicast mode.

When you have completed the fields on the General Settings tab, go to the Fencing tab.

Fencing

When certain problems occur on a server (for example, hardware problems or the loss of cluster network communications), and the server ceases to effectively coordinate and communicate with other servers in the cluster, HP Clustered File System must remove the server's access to filesystems to preserve data integrity. This step is called *fencing*.

The Fencing tab allows you to select the method that HP Clustered File System should use to remove access to the SAN.



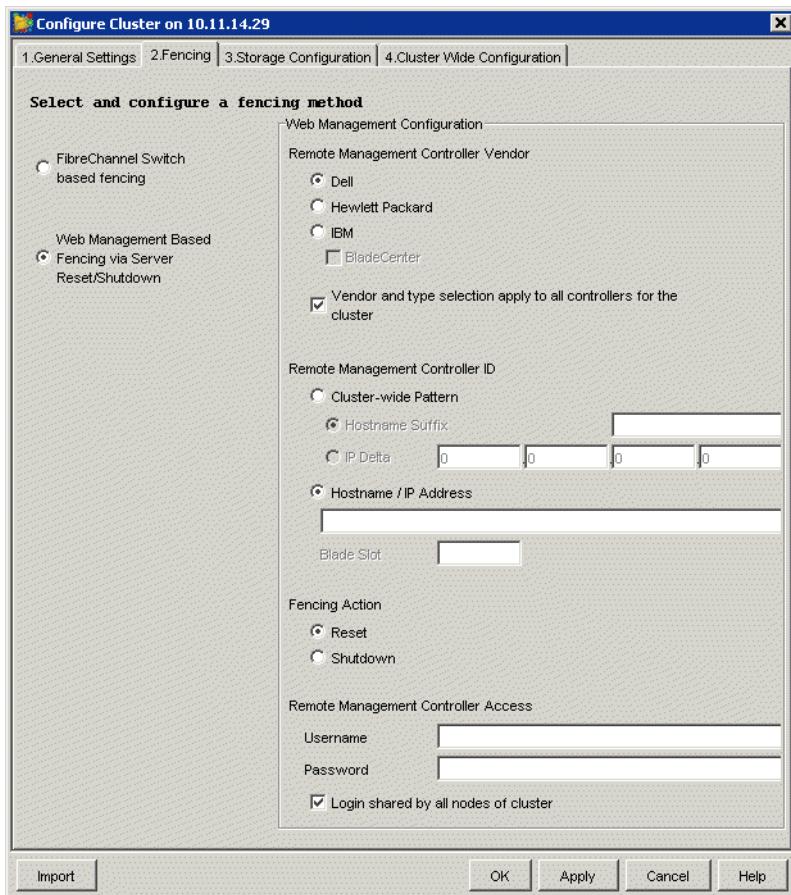
There are two fencing methods:

- **FibreChannel Switch-based fencing.** When a server needs to be fenced, HP Clustered File System will disable the server's access in the FibreChannel fabric. The server must be rebooted to regain access to the SAN. If you select this method, next go to the Storage Configuration tab and configure the FC switches. (See "Storage Configuration" on page 62.)
- **Web Management-Based Fencing via Server Reset/Shutdown.** HP Clustered File System uses remote management hardware on the server to remove its access to PSFS filesystems. (See the *HP Hardware and Software Compatibility Guide* on the HP web site for a list of supported servers and firmware.)

When you select Web Management-Based Fencing via Server Reset/Shutdown, the Fencing tab will refresh and ask you for configuration information. Enter the appropriate information for the Remote Management Controller associated with the server that you used to connect to the Management Console.

Later in this procedure the complete cluster configuration will be exported to the other servers that will be in the cluster. During the export, HP Clustered File System will, where possible, use the information that you specified on the Fencing tab to configure Web Management-Based Fencing on the other servers. If a particular entry on the Fencing tab did not apply to all servers, HP Clustered File System will ask for that information for each of the other servers.

NOTE: After HP Clustered File System is configured, you can verify that it has the correct information to fence each server. See "Test the Fencing Configuration" on page 68.



- 1. Remote Management Controller Vendor.** Select the vendor for your Remote Management Controllers. By default, the item “Vendor and type selection apply to all controllers for the cluster” is checked. Remove the checkmark from this item if your Remote Management Controllers are from different vendors.
- 2. Remote Management Controller ID.** Specify how HP Clustered File System should identify the Remote Management Controller associated with each server. Use one of the following methods.

- Select “Cluster-wide Pattern” and then specify the common naming scheme that you are using for the Remote Management Controllers (either a hostname suffix or an IP address delta).
- Enter the hostname or IP address for the Remote Management Controller associated with this cluster server. (If you unchecked “Vendor and type selection apply to all controllers for the cluster,” you will need to use this method.)

Hostname Suffix. Specify the common suffix to append to each server name to determine the associated Remote Management Controller name. For example, if your server names are server1 and server2 and their Remote Management Controllers are server1-iLO and server2-iLO, enter **-iLO** as the suffix.

IP Delta. Specify the delta to add to each server’s IP address to determine the IP addresses of the associated Remote Management Controllers. For example, if your servers are 1.255.200.12 and 1.255.200.15 and their Remote Management Controllers are 1.255.201.112 and 1.255.201.115, enter **0.0.1.100** as the delta.

Hostname/IP Address. Specify either the hostname or IP address for the Remote Management Controller associated with the server to which you are connected.

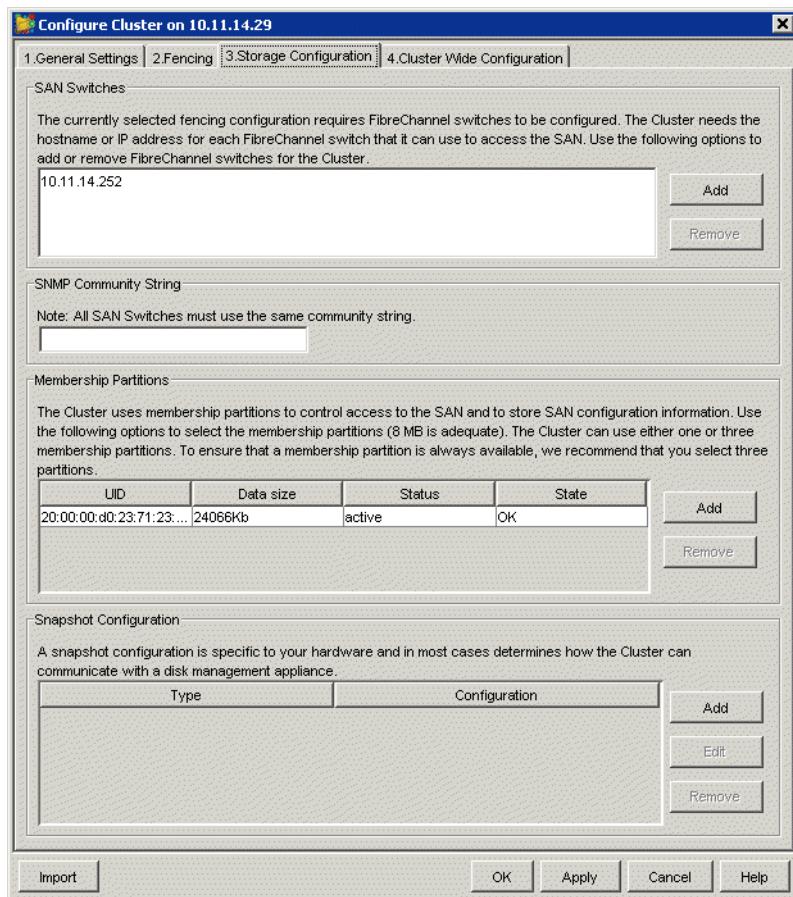
3. **Fencing Action.** When a server needs to be restricted from the SAN, HP Clustered File System can either power-cycle the server or shut it down. Select the method that you want to use.
4. **Remote Management Controller Access.** HP Clustered File System needs to log into a user account on the Remote Management Controller in order to fence the server. Specify the user name and password for the account that you want HP Clustered File System to use. The account must currently exist on the Remote Management Controller.

If all servers in the cluster will not be sharing the same login account, remove the checkmark from “Login shared by all nodes of cluster.”

When you have completed the Fencing tab, go to the Storage Configuration tab.

Storage Configuration

The Storage Configuration tab allows you to identify the FibreChannel switches included in the cluster and to select membership partitions, which HP Clustered File System uses to control access to the SAN.



1. **SAN Switches.** If you are using FibreChannel switch-based fencing, you will need to configure FC switches. If you are using Web Management-Based Fencing via server reset/shutdown, this step is optional; however, if the switches are configured the HP Clustered File System Management Console can display the switch ports used by the SAN. (The preceding window shows the text that appears for FibreChannel switch-based fencing.)

To configure SAN switches, you will need to specify the hostnames or IP addresses of the FibreChannel switches that are directly connected to the nodes in the cluster. Click Add, and then specify the hostname or IP address of the first FC switch. Repeat this procedure to specify the remaining FC switches, including cascading switches.



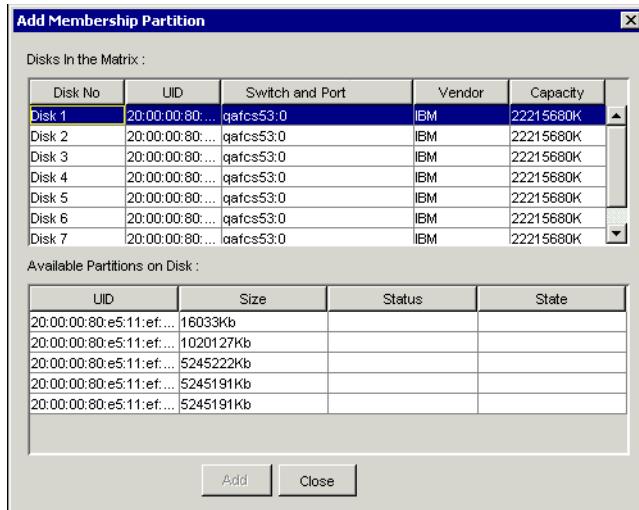
2. **SNMP Community String.** The default SNMP community string for HP Clustered File System is **private**. If you want to use a custom community string, enter the appropriate value here. The SNMP community string must be set to the same value on HP Clustered File System and on the SAN switches configured above.
3. **Membership Partitions.** HP Clustered File System uses a set of membership partitions to control access to the SAN and to store the device naming database, which includes the global device names that HP Clustered File System assigns to the SAN disks placed under its control.

You will need to select the LUNs or disk partitions that should be used as membership partitions.

NOTE: LUNs must already be partitioned as described in Chapter 2, under step “4. Create LUNs or Disk Partitions for Membership Partitions” on page 16.

HP Clustered File System can use either one or three membership partitions. **To ensure that a membership partition is always available, we strongly recommend that you select three partitions.**

To create a membership partition, click Add. The Add Membership Partition window then lists all of the disks or LUNs that it can access. Select the disk or LUN where you want to place the first membership partition.



All of the available partitions on that disk or LUN then appear in the bottom of the window. Select one of these partitions and click Add. (1 MB is adequate for a membership partition.) Repeat this procedure to select two more membership partitions. We recommend that the partitions be on different disks.

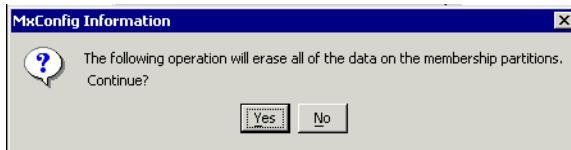
When selecting partitions for use as membership partitions, be sure that they do not contain any needed data. When the membership partitions are created, any existing data will be erased.

4. **Snapshot configuration.** HP Clustered File System provides support for taking hardware snapshots of PSFS filesystems. (The filesystems must be located on storage arrays supported for snapshots.) If you

want to use this capability, you will need to configure the snapshot method. Click Add, select the appropriate configuration module, and then supply the requested configuration information.

5. **Apply the configuration.** When you have completed your entries on the Storage Configuration tab, click Apply (at the bottom of the Cluster Configuration window).

The following message now appears. Click Yes to continue.

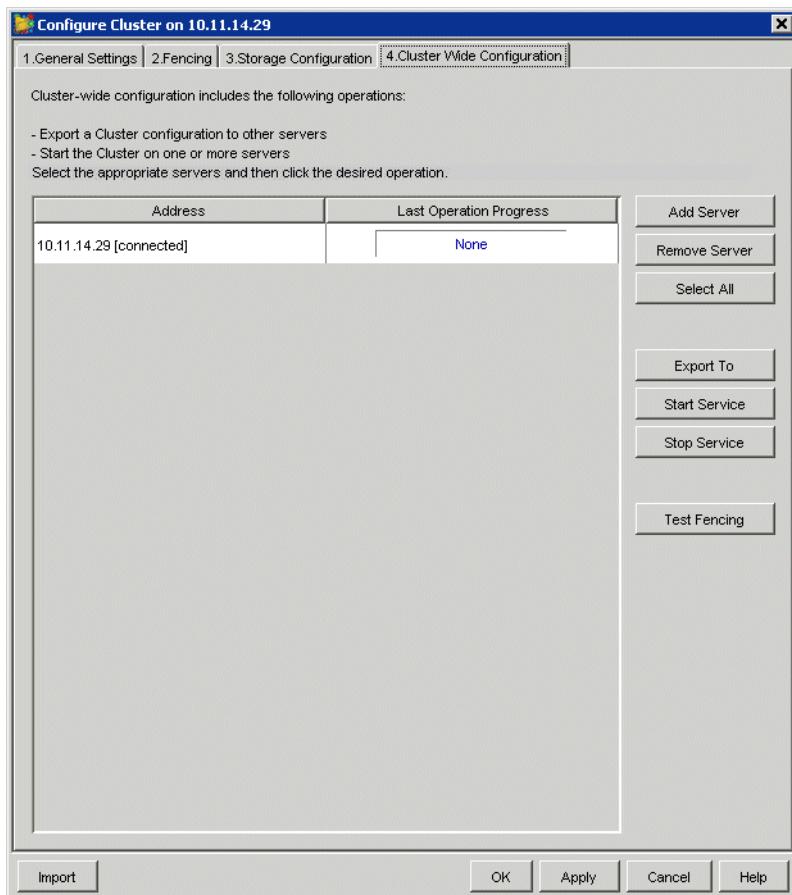


The configuration is then installed on the server that you are using to connect to the Management Console. You will then be asked whether you want to start the cluster on that server. Answer Yes.

The server is now configured and is running HP Clustered File System. Go to the Cluster-Wide Configuration tab.

Cluster Wide Configuration

This tab is used to export the cluster configuration to the other servers that will be in the cluster. It can also be used to start or stop HP Clustered File System on specific servers.



1. **Select the servers to be configured.** To specify the first server, click Add Server and type the hostname or IP address of the server on the Input form. Then click OK to add the server to the Address column.



Repeat this procedure to add the remaining servers to the Address column.

2. **Export the configuration.** Click Select All to select all of the servers in the Address column. Then click Export. The Last Operation Progress column will display status messages as the configuration is exported to each server.

If you are using Web Management-Based Fencing, you may be asked for additional information about each server. Based on the entries you made on the Fencing tab, the information can include any or all of the following: the vendor/type for the Remote Management Controller, the IP address/hostname of the Remote Management Controller, the username and password for the Remote Management Controller.

3. **Start HP Clustered File System on the remaining servers.** The servers to which you exported the configuration will still be selected in the Address column. Click Start Service to start HP Clustered File System on those servers. As each server is started, a status message will appear in the Last Operation Progress column.

When HP Clustered File System is running on all of the servers, you can close the Server Configuration window.

See the *HP StorageWorks Clustered File System Administration Guide* or the Management Console online help for information about configuring the cluster components.

Test the Fencing Configuration

The Test Fencing button on the Cluster Wide Configuration tab can be used to verify that the fencing configuration is correct for each server. This feature is particularly useful for Web Management-Based Fencing via Server Reset/Shutdown.

On the Cluster Wide Configuration tab, select one or more servers to test and click the Test Fencing button. (You cannot select the server being used to connect to the Management Console.)

HP Clustered File System then fences the servers in accordance with the method you specified on the Fencing tab. If the correct server is not fenced, you will need to check the configuration. For example, if you specified a pattern for the Remote Management Controller, the server may not be configured to use that pattern.

If an incorrect hostname or IP address was specified for the Remote Management Controller associated with the server, export the cluster configuration to that server again and then supply the correct information.

Index

A

add new server to cluster 50

C

change fencing method 52
CIFS file serving 48
client network connections 4, 6
cluster
 add new server 50
cluster nodes 4
clustered file system parameter
 set 53
clustered gateway
 setup 4
configuration
 information 1
 limits 1
configure cluster 54

D

disaster recovery 46

F

fencing configuration
 verifying 43
fencing method
 change 52

G

getting help v

H

HBA drivers 2
HP
 storage web site v

technical support v

I

iLO card network connections 6
Integrated Lights Out card 6
intra-cluster network connections 6

L

license file
 upgrade 51

M

Management Console
 install on other hosts 44
 Linux 44
 start 55
 uninstall 52
 Windows 45

O

other procedures 50

P

physical connectivity setup 4
post setup procedures 43
pre-planning steps 9

Q

quick restore 46
quota management 48

R

restore system state 46

S

set HP clustered file system
 parameter 53
setting up HP StorageWorks DL380-
 SL Clustered Gateway 4
setting up physical connectivity 4
setup checklist 10
setup procedure 9
software image overview 2
supported HBA drivers 2

T

technical support, HP v

U

uninstall Management Console 52
upgrade license file 51

V

verify fencing configuration 43